A note on the stability of the Swedish Phillips curve

Sune Karlsson¹ · Pär Österholm¹

Received: 26 June 2018 / Accepted: 15 July 2019 / Published online: 20 July 2019
© The Author(s) 2019

Abstract

We use Bayesian techniques to estimate bivariate VAR models for Swedish unemployment rate and inflation. Employing quarterly data from 1995Q1 to 2018Q3 and new tools for model selection, we compare models with time-varying parameters and/or stochastic volatility to specifications with constant parameters and/or covariance matrix. The evidence in favour of a stable dynamic relationship between the unemployment rate and inflation is mixed. Model selection based on marginal likelihood calculations indicates that the relation is time varying, whereas the use of the deviance information criterion suggests that it is constant over time; we do, however, note consistent evidence in favour of stochastic volatility. An out-of-sample forecast exercise is also conducted, but similarly provides mixed evidence regarding which model to favour. Importantly though, even if time-varying parameters are allowed for, our results do not suggest that the Phillips curve has been flatter in more recent years. This finding thereby questions the explanation that a flatter Phillips curve is the cause of the low inflation that Sweden has experienced in recent year.

Keywords Inflation · Unemployment · Time-varying parameters · Stochastic volatility

JEL Classification C11 · C32 · E32

We are grateful to Mikael Apel, Paola Di Casola, Henry Ohlsson, two anonymous referees, the editor, the associate editor and seminar participants at the Ministry of Finance (in Sweden), the 6th Swedish National Conference in Economics and the 10th Nordic Econometric Meeting for valuable comments. Rafael Barros De Rezende kindly provided data on the shadow rate and Joshua Chan generously shared Matlab code.

✉ Pär Österholm
par.osterholm@oru.se

Sune Karlsson
sune.karlsson@oru.se

¹ School of Business, Örebro University, 70182 Örebro, Sweden

1 Introduction

Inflation in Sweden has been stubbornly low over the last years. This development is similar to that in several other inflation-targeting countries where inflation has been moderate, and increasing slowly, despite historically low policy-interest rates and developments in the real economy which many argue should have generated a stronger inflationary pressure; see, for example, Jansson (2017) and Yellen (2017). One explanation for the low inflation which has been put forward is that the Phillips curve has become flatter, for example, due to globalisation and digitalisation.¹ Other explanations have, however, also been suggested. For instance, the oil price, aggregate demand, monetary policy and demographic trends have—just to mention a few variables—been put forward as having some importance; see, for example, Conti et al. (2017), Juselius and Takáts (2018) and Williamson (2018).

In this paper we contribute to the discussion regarding the properties of the Phillips curve by providing evidence based on Swedish data. Employing a Bayesian VAR (BVAR) framework, we estimate bivariate models using quarterly data on unemployment rate and inflation. The models are estimated under different assumptions concerning the dynamics and covariance matrix. As noted by a growing literature, time variation in both dynamics and volatility appears to be important features of macroeconomic relationships; see, for example, Cogley and Sargent (2005), Koop et al. (2009), Franta et al. (2014), Chan et al. (2016), Knotek and Zaman (2017) and Akram and Mumtaz (2019). We therefore estimate models with time-varying parameters and/or stochastic volatility and compare these models to specifications with constant parameters and/or covariance matrix. Four combinations are considered: (i) constant parameters and covariance matrix, (ii) time-varying parameters and constant covariance matrix, (iii) constant parameters and stochastic volatility and (iv) time-varying parameters and stochastic volatility. Relying on new tools for model selection developed by Chan and Eisenstat (2018), we formally assess which model is preferred by the data. This constitutes a step forward relative to the vast majority of previous related research. Since model selection is a non-trivial issue when models with time-varying parameters and stochastic volatility are involved, it has often simply been assumed that it is reasonable to employ a model with such features.² In this paper we instead evaluate this assumption using marginal likelihoods and the deviance information criterion (DIC) and can provide statistical evidence on the stability of the Swedish Phillips curve.

The focus in our analysis is put on bivariate models estimated using price inflation and the unemployment rate. While the original analysis conducted by Phillips (1958) was based on wage inflation, the more recent debate regarding the flattening of the

¹ For recent contributions to the debate regarding the potentially flattened Phillips curve, see Blanchard et al. (2015), Coibion and Gorodnichenko (2015) and Leduc and Wilson (2017). Another argument in favour of a flattened Phillips curve is the “missing disinflation” in the aftermath of the global financial crisis of 2008; see, for example, IMF (2013) for an overview. This missing disinflation has in and of itself been an issue of a fair amount of research; examples include Ball and Mazumder (2011), Matheson and Stavrev (2013) and Huang and Luo (2018).

² Exceptions do exist; see, for example, Koop et al. (2009) and Karlsson and Österholm (2018). In addition, studies such as D’Agostino et al. (2013) and Barnett et al. (2014) have relied on out-of-sample predictive criteria to distinguish between models.

Phillips curve has typically been focused on price inflation and we believe that this focus is reasonable given the many central banks—including Sveriges Riksbank—that target price inflation. However, we do not wish to go too far from Phillips' original idea and the intuitive bivariate setting; after all, we believe that there is value in assessing a reasonably clean and simple relation between inflation and the unemployment rate. That said, we also conduct analysis based on wage inflation and use different measures of resource utilization in order to assess the sensitivity of our findings from the benchmark specification (with price inflation and the unemployment rate). As part of this sensitivity analysis, the model is also augmented beyond the bivariate case with some key variables. Finally, we conduct an out-of-sample forecast exercise in order to further evaluate the relative merits of the different assumptions regarding the dynamics and covariance matrix.

Our main results suggest that the evidence in favour of a stable relation between the unemployment rate and inflation is mixed. Model selection based on marginal likelihoods indicates that the relation is time varying; if we instead rely on the DIC, the relation is judged to be constant over time. Importantly though, we do not—regardless of model specification—find any evidence suggesting that the Phillips curve has been flatter than usual during the last years of our sample and accordingly conclude that the low inflation in the last few years has not been caused by a changing slope of the Phillips curve. Our findings instead indicate that low trend inflation may have contributed to this development. Regarding the results from the sensitivity analysis, we argue that our main conclusion—that is, that there is no evidence of the Swedish Phillips curve being flatter during the last few years of the sample—is very robust. The results from the out-of-sample forecast exercise are similar to those from the within-sample analysis—that is, mixed. No model shows superior forecasting performance at all investigated horizons.

The rest of this paper is organised as follows: In Sect. 2, we describe the Bayesian VAR models which we employ. The results from our main empirical analysis relying on price inflation and the unemployment rate are presented and discussed in Sect. 3. In Sect. 4, we show results from various alternative specifications in order to assess how sensitive our main results are. An out-of-sample forecast exercise is conducted in Sect. 5. We provide an overall discussion of our results in Sect. 6. Finally, Sect. 7 concludes.

2 The Bayesian VAR models

We rely on BVARs for our analysis since the inflation equation of the BVAR can be seen as a “dynamic generalization of the Phillips curve” (King and Watson 1994, p. 172). While we pay special attention to the inflation equation, our analysis is mainly based on a bivariate system since important aspects of the dynamic relation between the variables otherwise could be lost.

Defining the vector of dependent variables as $\mathbf{y}_t = (u_t \pi_t)'$, where u_t is the unemployment rate and π_t is inflation, we specify the BVAR in its most general

form—that is, with time-varying parameters and stochastic volatility (“TVP-SV”)—in Eq. (1):

$$\mathbf{B}_{0t} \mathbf{y}_t = \boldsymbol{\gamma}_t + \mathbf{B}_{1t} \mathbf{y}_{t-1} + \dots + \mathbf{B}_{pt} \mathbf{y}_{t-p} + \boldsymbol{\varepsilon}_t \quad (1)$$

where $\mathbf{B}_{0t}$ is a 2×2 lower triangular matrix with ones on the diagonal, $\boldsymbol{\gamma}_t$ contains the time-varying intercepts and the matrices $\mathbf{B}_{1t}, \dots, \mathbf{B}_{pt}$ describe the dynamics. The vector of disturbances, $\boldsymbol{\varepsilon}_t$, follows $\boldsymbol{\varepsilon}_t \sim N(0, \boldsymbol{\Sigma}_t)$, where $\boldsymbol{\Sigma}_t = \text{diag}(\exp(h_{1t}), \exp(h_{2t}))$; the fact that the covariance matrix is diagonal and that the matrix $\mathbf{B}_{0t}$ premultiplies $\mathbf{y}_t$ means that the model is specified in a “structural” form.³ The lag length is set to $p = 4$. Collecting the free parameters of $\boldsymbol{\gamma}_t$ and $\mathbf{B}_{it}$ in the 19×1 parameter vector $\boldsymbol{\theta}_t$, we specify the processes for the time-varying parameters and log-volatilities as random walks:

$$\boldsymbol{\theta}_t = \boldsymbol{\theta}_{t-1} + \boldsymbol{\eta}_t \quad (2)$$

$$\mathbf{h}_t = \mathbf{h}_{t-1} + \boldsymbol{\zeta}_t \quad (3)$$

where $\boldsymbol{\eta}_t \sim N(0, \boldsymbol{\Sigma}_\theta)$ and $\boldsymbol{\zeta}_t \sim N(0, \boldsymbol{\Sigma}_h)$.

Having presented the most general version of the VAR model above, restrictions are then imposed when estimating the three alternative specifications, that is, (i) constant parameters and covariance matrix (“constant”), (ii) time-varying parameters and constant covariance matrix (“TVP”) and (iii) constant parameters and stochastic volatility (“SV”). When estimating a model with constant parameters, $\boldsymbol{\Sigma}_\theta$ is restricted to be zero. Similarly, $\boldsymbol{\Sigma}_h$ is restricted to be zero when a model with constant covariance matrix is estimated.

When it comes to the issue of model selection—that is, assessing which model is preferred by the data—we rely on marginal likelihoods as well as the DIC. In a Bayesian setting, the marginal likelihood is the appropriate measure of how well the model and prior agree with the data; the model with the highest marginal likelihood is the one preferred by the data. While the marginal likelihood is the prior expected value of the likelihood, the DIC (Spiegelhalter et al. 2002) combines a measure of fit, the posterior expectation of the log-likelihood, with a penalty term for the number of parameters and can thus be argued to be less sensitive to the choice of prior distributions.

As we rely on both marginal likelihoods and the DIC for model choice, we take care to ensure that the prior information is comparable across models. For the regression parameters, $\boldsymbol{\theta}$, we use an uninformative prior for the overall location together with a relatively tight prior on the evolution over time. Specifically, for models with constant parameters, the prior is $\boldsymbol{\theta} \sim N(0, \mathbf{V}_\theta)$ with $\mathbf{V}_\theta = 10\mathbf{I}$. For models with time-varying parameters, the same normal prior is used for the initial state, $\boldsymbol{\theta}_0$, and the evolution is governed by the variance of the increments. The diagonal elements of $\boldsymbol{\Sigma}_\theta$ are given

³ This interpretation of “structural”—used by, for example, Chan and Eisenstat (2018)—is different from the literature using “structural VARs” in order to identify “structural shocks” such as supply shocks or monetary policy shocks; see Bernanke (1986) or Blanchard and Quah (1989) for early contributions. We have no such ambitions in this paper.

inverse Gamma priors, $\sigma_{\theta i}^2 \sim iG(v_{\theta i}, S_{\theta i})$, with $v_{\theta i} = 5$ and prior means of 0.01 for the intercepts and 0.0001 for the other regression parameters. In expectation this yields a prior variance for the final state close to 11 for intercepts and close to 10 for the other parameters.

The specification of the prior for the error variances or volatilities requires more care as conditional conjugacy suggests inverse Gamma priors for the diagonal elements of Σ in the constant variance models and a normal prior for the initial state, $\mathbf{h}_0$, or a log-normal prior for the initial variance in the stochastic volatility models. As the first step, we tailor the prior to match the scale and variation of the data. Let s_i^2 be the residual variance from a univariate AR model. Employing a normal prior, $\mathbf{h}_0 \sim N(\mu_h, c\mathbf{I})$, we set $\mu_{h,i} = \ln s_i^2 - \frac{c}{2}$ to match the prior mean of $\exp(h_{it})$ with the residual variance. Similarly, for the constant variance case and inverse gamma priors, $\sigma_i^2 \sim iG(v_{0i}, S_{0i})$; for the diagonal elements of Σ , we set $S_{0i} = (v_{0i} - 1)s_i^2$. Finally, c and v_{0i} are selected so that the shapes of the prior distributions match while keeping them relatively uninformative, resulting in $c = 0.25$ and $v_{0i} = 5$. Finally, the prior for the diagonal elements of Σ_h is inverse Gamma, $\sigma_{hi}^2 \sim iG(v_{hi}, S_{hi})$, with $v_{hi} = 5$ and $S_{hi} = 0.04$ for a prior mean of the variance of 0.01.

While the models with time-varying parameters and/or stochastic volatility are inherently more flexible than the VAR with constant parameters and covariance matrix, the properties of the data implied by the priors are quite comparable across models. The prior means of the regression parameters are the same for the different models, and the increase of the prior variance with t is moderate for the models with time-varying parameters. While the prior expectation of the log-volatilities is constant over time with only a moderate increase in the variance, the log-normal form does imply that the volatility is increasing with t a priori. We have $E(\exp(h_{it})) = \exp(\mu_{h,i} + t\sigma_{hi}^2/2)$, and at the prior mean of σ_{hi}^2 , this implies that the prior expectation of the volatility at the end of the sample is 60% higher than the beginning of the sample. This is, however, more an issue of model specification than prior specification where most of the tractable model specifications that result in constant prior expectation of the volatilities imply decreasing prior expectation of the log-volatilities. For simplicity we have opted for the current, quite standard, specification of the time-varying volatilities.

For posterior inference, we rely on the Markov Chain Monte Carlo sampler employed by Chan and Eisenstat (2018). The sampler proceeds by sampling from the full conditional posteriors

$$p(\theta | \mathbf{y}, \mathbf{h}, \Sigma_\theta, \Sigma_h, \theta_0, \mathbf{h}_0) \quad (4)$$

$$p(\mathbf{h} | \mathbf{y}, \theta, \Sigma_\theta, \Sigma_h, \theta_0, \mathbf{h}_0) \quad (5)$$

$$p(\Sigma_\theta, \Sigma_h | \mathbf{y}, \theta, \mathbf{h}, \theta_0, \mathbf{h}_0) \quad (6)$$

$$p(\theta_0, \mathbf{h}_0 | \mathbf{y}, \theta, \mathbf{h}, \Sigma_\theta, \Sigma_h) \quad (7)$$

The full conditional posterior of $\theta = (\theta'_1, \dots, \theta'_T)'$ is normal both in models with constant and time-varying parameters. In models with stochastic volatility, $\mathbf{h}$ is sam-

pled using the auxiliary mixture sampler of Kim et al. (1998). The diagonal elements of Σ_θ and Σ_h are conditionally independent with inverse Gamma full conditional posteriors. Finally, θ_0 and h_0 are conditionally independent with normal full conditional posteriors. For the models with constant variance, the last three steps are replaced by an update of the diagonal elements of Σ which are conditionally independent with inverse Gamma full conditional posteriors.

The marginal likelihood and DIC calculations are carried out using the methods developed in Chan and Eisenstat (2018). Write the marginal likelihood for the TVP-SV model as

$$m(y) = \int p(y|\theta, h, \xi) p(\theta|\xi) p(h|\xi) p(\xi) d\theta dh d\xi \quad (8)$$

where ξ collects the parameters θ_0 , h_0 , Σ_θ and Σ_h of the state Eqs. (2) and (3). The integral with respect to θ is analytic, and importance sampling is used to integrate out h . Having integrated out the latent time-varying parameters and volatilities, this yields what Chan and Eisenstat term the integrated likelihood

$$p(y|\xi) = \int p(y|\theta, h, \xi) p(\theta|\xi) p(h|\xi) d\theta dh \quad (9)$$

The remaining integration with respect to the parameters in ξ is then carried out in an outer importance sampling loop to obtain $m(y) = \int p(y|\xi) p(\xi) d\xi$. For the TVP model the integrated likelihood calculation skips the integration with respect to h and in the outer loop ξ collects the parameters θ_0 , Σ_θ and $\text{diag}(\Sigma)$, while the SV model omits the integration with respect to θ in the integrated likelihood and ξ collects the parameters θ , h_0 and Σ_h . That is, the constant regression parameters are integrated out in the outer loop. For the constant parameter VAR, there are no latent parameters, the integrated likelihood coincides with the marginal likelihood, and the integration with respect to θ and $\text{diag}(\Sigma)$ in ξ is analytic.

The DIC is based on the deviance which, for the purpose of model comparison, can be defined as $D(\xi) = -2 \ln p(y|\xi)$. Model fit is measured by the posterior mean deviance

$$\overline{D(\xi)} = -2 \int \ln p(y|\xi) p(\xi|y) d\xi \quad (10)$$

That is, in contrast to the marginal likelihood, the integration is over the posterior distribution of ξ and not the prior. Model complexity is measured by the effective number of parameters,

$$p_d = \overline{D(\xi)} - D(\tilde{\xi}) \quad (11)$$

where $\tilde{\xi}$ is an estimate of ξ , here the posterior mean. Combining these yields the DIC as

$$\text{DIC} = \overline{D(\xi)} + p_d \quad (12)$$

Fig. 1 Unemployment rate and price inflation. *Note:* Variables measured in percent. Price inflation is given as the year-on-year change in the CPIF index. *Source:* Macrobond

The calculation of the integrated likelihood in (9) is done in the same way as for the marginal likelihood, and the integration with respect to ξ in (10) is carried out in an outer importance sampling loop.

3 Empirical findings: unemployment rate and price inflation

In this section, we use data on seasonally adjusted unemployment rate and CPIF⁴ inflation ranging from 1995Q1 to 2018Q3. The unemployment rate refers to the age group 16–64 years. CPIF inflation is calculated as $\pi_t^P = 100(P_t/P_{t-4} - 1)$, where P_t is the CPIF index at time t . We choose 1995Q1 as the starting point since it formally constitutes the start of Sweden's inflation-targeting regime and this gives the model with constant parameters and covariance matrix a fair chance to be considered a relevant modelling choice.⁵ Data are shown in Fig. 1.

We estimate the four different models presented in Sect. 2. Evaluation measures—in terms of log marginal likelihoods and DICs—are shown in Table 1. Judging by the log marginal likelihoods, the model with time-varying parameters and stochastic volatility is the preferred one. Based on this evidence, we would conclude that the Phillips curve has not been stable. However, Table 1 also shows that the DIC indicates that the

⁴ CPIF is the consumer price index with a fixed interest rate. As of September 2017, the Riksbank's inflation target is expressed in terms of the year-on-year percentage change in this index. Prior to that, the Riksbank was targeting the year-on-year change in the CPI. It can be noted that the only difference between the CPIF and the CPI is how the cost of housing is calculated (where the interest rate is assumed to be unchanged in the CPIF). While being less volatile than the CPI, the CPIF is not a measure of “core” inflation. Core inflation measures that the Riksbank relies upon include the CPIF excluding energy, TRIM85 and UND24; see, Johansson et al. (2018) for a further discussion.

⁵ The regime was declared in 1993 but was to start taking place from 1995; see Sveriges Riksbank (1993).

Table 1 Log marginal likelihood and DIC for the different BVAR specifications. Price inflation and unemployment rate data

Model	Log marginal likelihood	DIC
Constant	− 109.9	111.8
TVP	− 104.5	114.3
SV	− 107.7	110.1
TVP-SV	− 102.3	115.4

Based on bivariate BVARs with $\mathbf{y}_t = (u_t \pi_t^p)'$, where u_t is the unemployment rate and π_t^p is price inflation. The table gives the natural logarithm of the marginal likelihood of the models. “DIC” is the deviance information criterion. “Constant” is the BVAR model with constant covariance matrix and parameters. “TVP” is the BVAR model with constant covariance matrix and time-varying parameters. “SV” is the BVAR model with stochastic volatility and constant parameters. “TVP-SV” is the BVAR model with stochastic volatility and time-varying parameters

best model is that with constant parameters and stochastic volatility; perhaps somewhat surprisingly, it also suggests that the model with time-varying parameters and stochastic volatility is the worst. The method used for model selection obviously matters, and we conclude that this exercise has left us with contradictory results regarding whether there is time variation in the Swedish Phillips curve. It can be noted though that regardless of which selection tool that is used, a model with stochastic volatility comes out on top.

If the Phillips curve is stable over time—as the DIC suggests—the dynamic aspects that we are primarily interested in can quite easily be summarised. Looking at the preferred model—that is, the model with constant parameters and stochastic volatility—the first aspect of interest is the impulse-response function which describes how inflation reacts to a shock in the unemployment rate. This is shown in Fig. 2. As can be seen, a one-standard-deviation shock to the unemployment rate lowers—in line with our expectations—inflation for a number of quarters. The largest effect is found at the two-quarter horizon. It should be noted that the minor differences in the impulse-response function that can be found between different dates are only due to the fact that the size of the shock is time varying. The size of the shock—that is, its standard deviation—is illustrated in the impulse-response function which describes the effect that a shock to the unemployment rate has on the unemployment rate itself; see Fig. 11 in “Appendix”. As can be seen, there is limited variation in this standard deviation—it ranges from 0.20 to 0.25. The evidence for stochastic volatility in the model appears to be driven by the volatility of the shocks to the inflation equation which is shown in Fig. 12.⁶

The second aspect of the Phillips curve that is of primary interest is its slope. In line with other related research—see, for example, Knotek and Zaman (2017) and Karlsson and Österholm (2018)—we report the sum of the coefficients on the

⁶ For completeness, the impulse-response function which describes the effect that a shock to inflation has on the unemployment rate is shown in Fig. 13. Also this effects seems reasonable (in line with a cost-push shock).

Fig. 2 Impulse-response function: Effect of shocks to the unemployment rate on price inflation. Model with constant parameters and stochastic volatility. *Note:* Based on bivariate BVAR with $y_t = (u_t \pi_t^p)'$, where u_t is the unemployment rate and π_t^p is price inflation. Size of impulse is one standard deviation. Effect in percentage points on vertical axis. Horizon in quarters and dates on horizontal axes

lagged unemployment rate in the inflation equation of the model's "reduced-form" specification. The reduced form of the VAR in Eq. (1) is given by

$$y_t = \delta_t + A_{1t}y_{t-1} + \cdots + A_{pt}y_{t-p} + e_t \quad (13)$$

where $\delta_t = B_{0t}^{-1} \gamma_t$, $A_{it} = B_{0t}^{-1} B_{it}$ and $e_t = B_{0t}^{-1} \varepsilon_t$. It can be noted that for the model with constant parameters and stochastic volatility, there is no time variation in γ_t , B_{0t} or B_{it} ; the slope of the Phillips curve is accordingly constant during the investigated sample, and it is estimated to be -0.06 . This is a flatter Phillips curve than what Karlsson and Österholm (2018) found for most periods for the USA (comparing point estimates) but very close to the slope that Knotek and Zaman (2017) found for the US Phillips curve in the two time points that they focused on, namely 1999Q3 and 2017Q3.

As shown in Table 1 though, the marginal likelihoods indicated that the model with time-varying parameters and stochastic volatility was the preferred one. We therefore next turn our attention to the properties of the Phillips curve based on this model. The impulse-response function describing the effect that a shock to the unemployment rate has on inflation is given in Fig. 3⁷; the other impulse-response functions are shown in Figs. 14, 15 and 16 in "Appendix". Also in this model, an unexpectedly high unemployment rate tends to decrease inflation at short horizons (apart from the three and four quarter horizons, the latter only in the beginning of the sample). The impulse-response function is reasonably stable over time despite the fact that the

⁷ The shock size is one standard deviation. This varies between 0.11 and 0.17 percentage points; see Fig. 14 in "Appendix".

Fig. 3 Impulse-response function: Effect of shocks to the unemployment rate on price inflation. Model with time-varying parameters and stochastic volatility. *Note:* Based on bivariate BVAR with $y_t = (u_t \pi_t^p)'$, where u_t is the unemployment rate and π_t^p is price inflation. Size of impulse is one standard deviation. Effect in percentage points on vertical axis. Horizon in quarters and dates on horizontal axes

estimated parameters of the model show time variation; see Fig. 4 for the inflation equation and Fig. 17 in “Appendix” for the unemployment rate equation.⁸ It can also be noted that the time variation in the equation for inflation is similar in magnitude to that in the equation for the unemployment rate. We thus conclude that the time variation suggested by the marginal likelihoods is not due to time variation in the equation for the unemployment rate alone.

Turning to the slope of the Phillips curve, this is plotted in Fig. 5. Taking the 68% credible interval into account, one might claim that the slope of the Phillips curve has not changed dramatically during the sample. However, looking at the point estimate, the story is somewhat different as it ranges from -0.10 to -0.35 . Interestingly though, the slope has not been unusually close to zero between 2011 and 2016 indicating that Sweden’s low inflation in this period cannot be explained by a flat Phillips curve.

The final aspect of the evolving properties of the model with time-varying parameters and stochastic volatility which we consider is a concept called “trend inflation”.

⁸ Looking at Fig. 4, it is also worth noting that the sum of the coefficients on lagged inflation is close to zero. Other things equal, a sum closer to zero (rather than closer to unity) means that inflation’s deviation from what we refer to as “trend inflation” below will be more short-lived. One explanation for this finding is the moderate persistence that Swedish CPIF inflation shows in general; the first-order autocorrelation is 0.77. This can be contrasted with US data where inflation is often modelled as having a unit root; see, for example, Stock and Watson (2009). Given that the model with time-varying parameters also can remove a persistent component from the series through the time-varying intercept, we do not consider our finding of a sum close to zero on the coefficients on lagged inflation to be very surprising. It is also in line with the findings of modest inflation gap persistence in the USA found by Cogley and Sbordone (2009) for the period 1984–2003—a period during which monetary policy generally is considered to have been focused on inflation stabilization. In contrast, Cogley and Sbordone show that between 1960 and 1983, inflation gap persistence was substantially higher. Further discussions concerning inflation gap persistence in the USA can be found in, for example, Benati and Surico (2008) and Cogley et al. (2009).

Fig. 4 Estimated coefficients of the inflation equation from the reduced form BVAR in Eq. (13). Price inflation and unemployment rate data. Model with time-varying parameters and stochastic volatility. *Note:* Based on bivariate BVAR with $y_t = (u_t \pi_t^p)'$, where u_t is the unemployment rate and π_t^p is price inflation

Fig. 5 Estimated “slope” of the Phillips curve. Price inflation and unemployment rate data. Model with time-varying parameters and stochastic volatility. *Note:* Based on bivariate BVAR with $y_t = (u_t \pi_t^p)'$, where u_t is the unemployment rate and π_t^p is price inflation. Coloured band is 68% equal-tail credible interval

In line with the terminology of Faust and Wright (2013) and Clark and Doh (2014), we define trend inflation as the value to which the inflation forecasts from the BVAR model converge.⁹ In order to define the concept algebraically, we first rewrite the reduced form VAR in Eq. (13) in companion form as

$$\tilde{y}_t = \tilde{\delta}_t + A_t \tilde{y}_{t-1} + \tilde{e}_t \quad (14)$$

⁹ The concept of trend inflation used in this paper accordingly has the same interpretation as what Cogley and Sargent (2005) refer to as “core inflation”.

Fig. 6 Estimated “trend inflation”. Price inflation and unemployment rate data. Model with time-varying parameters and stochastic volatility. *Note:* Based on bivariate BVAR with $y_t = (u_t \pi_t^p)'$, where u_t is the unemployment rate and π_t^p is price inflation. Coloured band is 68% equal-tail credible interval

with

$$\tilde{y}_t = (y'_t, y'_{t-1}, \dots, y'_{t-p+1})', \quad (15)$$

$$\tilde{\delta}_t = (\delta'_t, 0', \dots, 0')', \quad (16)$$

$$A_t = \begin{pmatrix} A_{1t} & A_{2t} & \cdots & A_{p-1,t} & A_{pt} \\ I & 0 & \cdots & & 0 \\ 0 & \ddots & & & \\ \vdots & & \ddots & & \vdots \\ 0 & \cdots & 0 & I & 0 \end{pmatrix} \quad (17)$$

and

$$\tilde{e}_t = (e'_t, 0', \dots, 0')' \quad (18)$$

Trend inflation is given as the second element of $\phi_t = (I - A_t)^{-1} \tilde{\delta}_t$. Estimated median trend inflation at each point in time is shown in Fig. 6 and is interesting to discuss from a monetary policy perspective.

Around 2010–2014, the Riksbank was “leaning against the wind” in order to dissuade households from taking on housing-related debt. This contributed to the low inflation outcomes during this period and sparked an intense debate regarding whether

the Riksbank's policy was expansive enough; see, for example, Svensson (2014). As can be seen, trend inflation fell while the Riksbank was leaning against the wind. By 2014 trend inflation was between 0.6 and 1%. This coincides with when the Riksbank abandoned the policy of leaning against the wind and declared that its focus was on achieving the inflation target—partly because they were concerned about de-anchoring of inflation expectations.¹⁰ The model results hence give some support for the Riksbank's decision to change its policy in 2014 since it suggests that the low inflation outcomes had become problematic as its long-run inflation forecasts were not consistent with the inflation target. In order to fight low inflation and low inflation expectations, the Riksbank lowered the repo rate to zero in October 2014 and kept it at -0.5% between February 2016 and December 2018.¹¹ After the change in the Riksbank's policy, inflation has slowly but surely increased and trend inflation with it; this development has occurred despite the weak development in oil prices that took place during part of this period (see Fig. 25 in “Appendix”), suggesting that a monetary policy focused on achieving the inflation target has been effective in moving both inflation outcomes and expectations towards the target.

Summarising the results so far, we note that there is mixed evidence on whether the Phillips curve is stable or not. However, even if we allow for the possibility that it has been unstable, it does not show signs of having been flatter than usual in the last few years of the analysed sample.

4 Empirical findings: sensitivity analysis

In order to assess the robustness of our findings in the previous section, we next conduct additional analysis that covers some relevant aspects given our empirical setting. First, we vary some of the prior settings of the models. Second, data on wage inflation are used instead of price inflation. Third, we replace the unemployment rate with alternative measures of resource utilization. Finally, we move beyond the bivariate framework and estimate models which include additional variables that could be of particular interest when studying Swedish inflation.

4.1 Priors

We believe the choice of priors described in Sect. 2—which were used when estimating the models underlying the results presented in Sect. 3—to be uncontroversial. The prior settings are, for example, close to the ones used by Chan and Eisenstat (2018). Still, given the stated purpose of facilitating model comparison, we have to some extent deviated from what can be seen as “common practice” with different model classes. While Minnesota style priors are common with constant parameter VARs, we do not use this type of prior as the increased prior shrinkage towards zero with increasing lag

¹⁰ Signs of de-anchored inflation expectations at this point in time include the five-year inflation expectations in Sweden's most prominent survey on inflation expectations—the TNS Prospera survey—which had fallen from 2.4% in late 2011 to 1.8% in mid-2014.

¹¹ In December 2018, the repo rate was increased to -0.25% . This is not included in our sample though.

Table 2 Log marginal likelihood for the different BVAR specifications. Price inflation and unemployment rate data

Model	Benchmark	$V_\theta = 20I$	$V_\theta = 5I$	$S_{\theta i} = 0.08, 0.0008$	$S_{\theta i} = 0.02, 0.0002$	$S_{hi} = 0.08$	$S_{hi} = 0.02$
Constant	− 109.9	− 116.3	− 103.5	− 109.9	− 109.9	− 109.9	− 109.9
TVP	− 104.5	− 110.4	− 98.8	− 109.9	− 103.9	− 104.5	− 104.5
SV	− 107.7	− 114.2	− 101.4	− 107.7	− 107.7	− 107.8	− 108.0
TVP-SV	− 102.3	− 108.3	− 96.9	− 106.7	− 102.6	− 102.4	− 102.8

Based on bivariate BVARs with $y_t = (u_t, \pi_t^p)'$, where u_t is the unemployment rate and π_t^p is price inflation. Table gives the natural logarithm of the marginal likelihood of the models. “Constant” is the BVAR model with constant covariance matrix and parameters. “TVP” is the BVAR model with constant covariance matrix and time-varying parameters. “SV” is the BVAR model with stochastic volatility and constant parameters. “TVP-SV” is the BVAR model with stochastic volatility and time-varying parameters

length to some extent is in conflict with the nature of time-varying parameter models. Similarly, while training sample priors are common in the literature using VARs with time-varying parameters, we do not use this type of prior since the monetary policy regime was quite different in the period before our sample and there are clear level breaks in both inflation and unemployment just prior to the start of our sample. Given this, a training sample prior would put constant parameter models at a disadvantage, which we want to avoid. It is nevertheless of interest to see how sensitive our results are with respect to the priors. More specifically, we vary the prior settings by halving and doubling key prior parameters one at a time. Specifically, we vary the prior variance of the regression parameters (or their initial state), V_θ , and the scale parameters $S_{\theta i}$ and S_{hi} governing the prior mean of the innovation variances in the state Eqs. (2) and (3).

In Tables 2 and 3, we present log marginal likelihoods and DICs for the four specifications under consideration using six different variations in the priors; the benchmark specification ($V_\theta = 10I$, $S_{\theta i} = 0.04$ for constants and 0.0004 for the other regression parameters and $S_{hi} = 0.04$) used in Sect. 3 is given in the result column furthest to the left. As can be seen from the two tables, the results with respect to model selection are extremely robust: Based on marginal likelihoods, the model with time-varying parameters and stochastic volatility is always the preferred one. Similarly, when the DIC is employed, the model with constant parameters and stochastic volatility is always judged as being superior to the others.

Based on the results above, we conclude that our results regarding model selection are not particularly sensitive with respect to our choice of prior parameters.

4.2 Wage inflation

We next turn our attention to the properties of the Phillips curve when it is estimated using data on wage inflation instead of price inflation. This has the benefit of taking our analysis in the direction of Phillips’ (1958) original observation which—as stated above—was based on wage inflation. But perhaps more importantly, it also means that we can scrutinise what may be described as a key relationship through which central

Table 3 DIC for the different BVAR specifications. Price inflation and unemployment rate data

Model	Benchmark	$V_\theta = 20I$	$V_\theta = 5I$	$S_{\theta i} = 0.08, 0.0008$	$S_{\theta i} = 0.02, 0.0002$	$S_{hi} = 0.08$	$S_{hi} = 0.02$
Constant	111.8	111.9	111.6	111.8	111.8	111.8	111.8
TVP	114.3	114.3	114.7	124.6	111.4	114.3	114.3
SV	110.1	110.2	110.0	110.2	110.2	110.4	110.5
TVP-SV	115.4	115.4	115.8	125.1	112.5	116.1	115.3

Based on bivariate BVARs with $y_t = (u_t \pi_t^p)'$, where u_t is the unemployment rate and π_t^p is price inflation. Table gives the DIC—that is, the deviance information criterion—of the models. “Constant” is the BVAR model with constant covariance matrix and parameters. “TVP” is the BVAR model with constant covariance matrix and time-varying parameters. “SV” is the BVAR model with stochastic volatility and constant parameters. “TVP-SV” is the BVAR model with stochastic volatility and time-varying parameters

Fig. 7 Unemployment rate and wage inflation. *Note:* Variables measured in percent. Wage inflation is given as the year-on-year change in the hourly wage. *Source:* Macrobond and National Institute of Economic Research

banks assume that their goals are achieved: Expansive monetary policy is assumed to stimulate the real economy, thereby lowering the unemployment rate; this tighter labour market pushes up wage growth (wage inflation) which in turn increases price inflation. Assessing the stability of the relation between the unemployment rate and wage inflation is thus of considerable interest.

Price inflation in the bivariate BVARs is accordingly replaced with wage inflation; this is calculated as $\pi_t^w = 100(W_t/W_{t-4} - 1)$, where W_t is the hourly wage at time t .¹² Our vector of dependent variables thus becomes $y_t = (u_t \pi_t^w)'$. As before, our sample ranges from 1995Q1 to 2018Q3. Data are shown in Fig. 7.

¹² The measure is based on the Swedish National Mediation Office's data and reflects the entire economy (that is, businesses and public sector).

Table 4 Log marginal likelihood and DIC for the different BVAR specifications. Wage inflation and unemployment rate data

Model	Log marginal likelihood	DIC
Constant	− 97.5	84.5
TVP	− 89.6	82.7
SV	− 91.4	79.8
TVP-SV	− 85.9	81.9

Based on bivariate BVARs with $y_t = (u_t \pi_t^w)'$, where u_t is the unemployment rate and π_t^w is wage inflation. Table gives the natural logarithm of the marginal likelihood of the models. “DIC” is the deviance information criterion. “Constant” is the BVAR model with constant covariance matrix and parameters. “TVP” is the BVAR model with constant covariance matrix and time-varying parameters. “SV” is the BVAR model with stochastic volatility and constant parameters. “TVP-SV” is the BVAR model with stochastic volatility and time-varying parameters

As was the case when the models were estimated using price-inflation data, we find that the model with time-varying parameters and stochastic volatility is the preferred one when looking at marginal likelihoods. However, the DICs once again conclude that the model with constant parameters and stochastic volatility is most suitable (Table 4).

The model with constant parameters and stochastic volatility obviously implies that the Phillips curve has been stable during the studied period and it is therefore not particularly interesting to study it in detail. In Fig. 18 in “Appendix”, we nevertheless show the impulse-response function describing the effect that a shock to the unemployment rate has on wage inflation.¹³ It can be noted though that while the effect in the short run is the expected—that is, wage inflation decreases—the slope of the Phillips curve for this model is 0.00. So while the model suggests that the Phillips curve has been stable, it can be questioned whether there actually exists a Phillips curve in this case.

If we instead turn to the preferred model according to marginal likelihoods, Fig. 8 shows the effect that a shock to the unemployment rate has on wage inflation.¹⁴ In this model with time-varying parameters and stochastic volatility, a shock to the unemployment rate tends to decrease wage inflation; the maximum effect of approximately −0.06 is reached after five quarters. And similarly to the case when using price inflation, the impulse-response function is reasonably stable over time.

In Fig. 9, the slope of the Phillips curve is presented. As can be seen, this shows a tendency to have become steeper over time, though one should of course recall that there is a fair bit of uncertainty associated with the estimates as indicated by the 68% credible interval. Importantly though, there is no evidence of the Phillips curve having flattened during the last few years of the sample.

Finally, in Fig. 10 we show the model’s estimate of trend inflation. This indicates that there have been noticeable differences in trend inflation over time. While it hovered between 4 and 5% from 1996 to 2002, it has varied around much more modest levels

¹³ The shock size is one standard deviation, which varies between 0.19 and 0.25.

¹⁴ The shock size is one standard deviation, which varies between 0.10 and 0.16.

Fig. 8 Impulse-response function: Effect of shocks to the unemployment rate on wage inflation. Model with time-varying parameters and stochastic volatility. *Note:* Based on bivariate BVAR with $y_t = (u_t \pi_t^w)'$, where u_t is the unemployment rate and π_t^w is wage inflation. Size of impulse is one standard deviation. Effect in percentage points on vertical axis. Horizon in quarters and dates on horizontal axes

Fig. 9 Estimated “slope” of the Phillips curve. Wage inflation and unemployment rate data. Model with time-varying parameters and stochastic volatility. *Note:* Based on bivariate BVAR with $y_t = (u_t \pi_t^w)'$, where u_t is the unemployment rate and π_t^w is wage inflation. Coloured band is 68% equal-tail credible interval

after the global financial crisis; in fact, trend inflation has not been more than 3% since 2009. Given that actual wage inflation has been low since 2010, this is not particularly surprising though.

Fig. 10 Estimated “trend inflation”. Wage inflation and unemployment rate data. Model with time-varying parameters and stochastic volatility. *Note:* Based on bivariate BVAR with $y_t = (u_t \pi_t^w)'$, where u_t is the unemployment rate and π_t^w is wage inflation. Coloured band is 68% equal-tail credible interval

Overall, the results based on wage inflation are largely in line with what we found using price inflation: The same two types of models were preferred by the data, and the evidence regarding the stability of the Phillips curve is mixed. Also, there is no support to a claim of a Phillips curve that has been flatter than usual in the last few years. However, when using wage inflation data, the model with constant parameters and stochastic volatility finds little support for a Phillips curve as the slope is estimated to be zero.¹⁵

4.3 Other measures of resource utilization

Given Phillips’ (1958) original observation, the unemployment rate is the most natural choice of variable to relate to inflation when conducting the kind of analysis that we do in this paper. However, during the process of becoming popular, the meaning of the term “Phillips curve” has become somewhat blurred. It now comes in many different forms and using a range of different variables.¹⁶ This fact is reflected in a description of the Phillips curve provided by Clark and McCracken (2006, p. 1127), who mean that it broadly can be defined as “... a model relating inflation to the unemployment rate, output gap, or capacity utilization”.

¹⁵ On the other hand, when using price inflation data in Sect. 3, the 68% credible interval for the slope was all in the negative range, from -0.09 to -0.02 .

¹⁶ One of the most important versions in modern macroeconomics is the so-called New Keynesian Phillips curve; see, for example, Galí and Gertler (1999) and Rudd and Whelan (2005). Recent empirical literature taking a less structural approach and relying on various techniques from time series econometrics includes Svensson (2015), Busetti and Caivano (2016), Chan et al. (2016) and Knotek and Zaman (2017).

Seeing that different measures of capacity utilization are considered in the empirical literature related to the Phillips curve, we next investigate the sensitivity of our results with respect to this aspect.¹⁷ We do this by estimating bivariate models with price inflation but using either the unemployment gap or the output gap instead of the unemployment rate. That is, we define the vector of dependent variables in Eq. (1) as $y_t = (u_t^{\text{gap}} \pi_t^p)'$ or $y_t = (g_t^{\text{gap}} \pi_t^p)'$ where u_t^{gap} is the unemployment gap and g_t^{gap} is the output gap.

The unemployment gap is defined as the actual unemployment rate minus the equilibrium (or “natural”) unemployment rate. If the equilibrium rate is constant, then it is irrelevant whether the unemployment gap or the unemployment rate is used. However, there are both theoretical and empirical arguments suggesting that the equilibrium rate is time varying; see, for example, Hall (2003). We account for this fact by using the unemployment gap as calculated by the National Institute of Economic Research (NIER).¹⁸ This gap is based on the NIER’s best estimate of a time-varying equilibrium unemployment rate in Sweden.¹⁹ Data on both the unemployment gap and the equilibrium unemployment rate are given in Fig. 19 in “Appendix”.²⁰ As can be seen, the equilibrium rate is assumed to move very moderately over time.

Turning to the output gap, this is also based on data from the NIER. It is calculated as the percentage deviation between actual GDP and the NIER’s best estimate of potential GDP. Data on the output gap can also be found in Fig. 19 in “Appendix”.

Results from the model selection exercise are given in Table 5. As can be seen, we find exactly the same results as before, regardless of capacity utilization measure: According to marginal likelihoods, the model with time-varying parameters and stochastic volatility is preferred whereas the DIC chooses the model with constant parameters and stochastic volatility.

Focusing on the slope of the Phillips curve in the models with constant parameters and stochastic volatility, we can note that the model with the unemployment gap has a slope of -0.07 , which is very close to what we found in our main analysis in Sect. 3 (that is, -0.06). The model estimated with output gap data has a positive slope as expected, namely 0.03 . Turning to the models with time-varying parameters and stochastic volatility, the slope of the Phillips curve when the model was estimated using the unemployment gap is shown in Fig. 20 in “Appendix”. This suggests that

¹⁷ Examples include broad unemployment, underemployment, the unemployment gap, the output gap, a labour market conditions indicator gap and labour underutilization; see Riggi and Venditti (2015), Bulligan et al. (2017), Bell and Blanchflower (2018) and Conti and Gigante (2018) for some recent examples. To some extent the use of alternative measures has been driven by a failure to find support for a Phillips curve when using specifications based on a “traditional” unemployment rate, unlike what we do in this paper. It is unclear why the unemployment rate appears to carry informational value for inflation in Sweden but not in several other countries. Given the complexity of the issue though, we refrain from suggesting explanations rather than providing what would have to be speculations.

¹⁸ The NIER is a Swedish government authority under the Ministry of Finance which does analysis and forecasting. <https://www.konj.se/>.

¹⁹ See, for example, National Institute of Economic Research (2017) for a discussion.

²⁰ It can be noted that the unemployment gap is calculated based on the unemployment rate and equilibrium unemployment rate for the age group 15–74 years. The reason why we do not use the same age group as in the main analysis (16–64 years) is that the NIER does not calculate an equilibrium unemployment rate for this age group. We believe that the age group 16–64 years is the most relevant though when looking at the Swedish labour market.

Table 5 Log marginal likelihood and DIC for the different BVAR specifications. Price inflation data

Model	Unemployment gap		Output gap	
	Log marginal likelihood	DIC	Log marginal likelihood	DIC
Constant	− 110.2	112.7	− 224.5	344.7
TVP	− 104.9	116.9	− 219.6	348.6
SV	− 107.2	109.3	− 218.0	333.4
TVP-SV	− 102.4	116.6	− 215.9	344.1

Based on bivariate BVARs with $y_t = (u_t^{\text{gap}} \pi_t^{\text{p}})'$ and $y_t = (g_t^{\text{gap}} \pi_t^{\text{p}})'$, where u_t^{gap} is the unemployment gap, g_t^{gap} is the output gap and π_t^{p} is price inflation. Table gives the natural logarithm of the marginal likelihood of the models. “DIC” is the deviance information criterion. “Constant” is the BVAR model with constant covariance matrix and parameters. “TVP” is the BVAR model with constant covariance matrix and time-varying parameters. “SV” is the BVAR model with stochastic volatility and constant parameters. “TVP-SV” is the BVAR model with stochastic volatility and time-varying parameters

our previous findings are robust; there are no signs that the Phillips curve has been flatter than usual the last few years. The same conclusion is drawn also when using output gap data—see Fig. 21 in “Appendix”. It can be noted though that the Phillips curve has the wrong slope in the beginning of the sample. However, given the dramatic developments in the Swedish economy after the financial crisis of the early 1990s, it is perhaps not completely surprising that some counterintuitive results emerge at that point in time.

Summing up the results, we note that using the unemployment gap or output gap as measures of resource utilization yields results that are quite similar to when the unemployment rate is used. We find no evidence of a flattening Phillips curve during the last years of the sample.

4.4 Additional factors

The models that we have relied upon so far have exclusively been bivariate. While we believe that this has many benefits, it is obviously the case that other factors than the unemployment rate might affect Swedish inflation. We therefore want to investigate the sensitivity of our findings to the use of somewhat larger models, keeping in mind that the purpose of this paper clearly is not to find the best analysis or forecasting model for Swedish inflation. In this subsection, we consider two aspects that appear particularly relevant for inflation, namely foreign variables and Swedish interest rates.

We first augment the model with the repo rate, i_t^r , which is the policy rate of the Riksbank. Data are shown in Fig. 22 in “Appendix”. This makes the vector of dependent variables $y_t = (u_t \pi_t^{\text{p}} i_t^r)'$. Trivariate models based on the unemployment rate, inflation and an interest rate are commonly employed tools in the related empirical literature using Bayesian VARs—see, for example, Cogley and Sargent (2005) and Primiceri (2005)—and accordingly seem like a natural extension in this case.

Results from the model selection exercise are shown in Table 6, and as can be seen, the model with time-varying parameters and stochastic volatility is chosen when we

Table 6 Log marginal likelihood and DIC for the different trivariate BVAR specifications using Swedish interest rates. Price inflation and unemployment rate data

Model	Repo rate		Shadow rate	
	Log marginal likelihood	DIC	Log marginal likelihood	DIC
Constant	− 181.5	119.0	− 181.7	126.0
TVP	− 177.8	148.1	− 178.5	153.7
SV	− 178.8	120.1	− 178.4	124.4
TVP-SV	− 173.3	148.5	− 175.2	153.4

Based on trivariate BVARs with $y_t = (u_t \pi_t^p i_t^r)'$ and $y_t = (u_t \pi_t^p i_t^s)'$, where u_t is the unemployment rate, π_t^p is price inflation, i_t^r is the repo rate and i_t^s is the shadow rate of De Rezende and Ristinieni (2018) based on a three-factor model. Table gives the natural logarithm of the marginal likelihood of the models. “DIC” is the deviance information criterion. “Constant” is the BVAR model with constant covariance matrix and parameters. “TVP” is the BVAR model with constant covariance matrix and time-varying parameters. “SV” is the BVAR model with stochastic volatility and constant parameters. “TVP-SV” is the BVAR model with stochastic volatility and time-varying parameters

rely on marginal likelihoods; the DIC, however, once again selects the model with constant parameters and stochastic volatility. The slope of the Phillips curve—where we still measure this as the sum of the coefficients on lagged unemployment in the inflation equation—for the latter model is -0.11 . For the model with time-varying parameters and stochastic volatility, the slope is shown in Fig. 23 in “Appendix”. Comparing it to the slope from the benchmark model (in Fig. 5), it can be seen that adding the repo rate to the system does not substantially affect the estimate.

However, in the sample employed here, traditional monetary policy has been restricted by the “zero lower bound” which lead the Riksbank to also conduct unconventional monetary policy, for example, by buying government bonds in large quantities.²¹ In order to take both conventional and unconventional monetary policy into account, we also estimate a model with a shadow rate (instead of the repo rate)²²; we accordingly define $y_t = (u_t \pi_t^p i_t^s)'$, where i_t^s is the shadow rate.²³ Results from the estimation are shown in Table 6.

As can be seen, we find results that echo our previous findings: Marginal likelihoods favour the model with time-varying parameters and stochastic volatility, and the DIC selects the model with constant parameters and stochastic volatility. The slope of the model with time-varying parameters and stochastic volatility is shown in Fig. 24 in “Appendix” and shows a striking resemblance to the results based on the model using the repo rate (in Fig. 23). The model with constant parameters and stochastic volatility has an estimated slope of -0.13 , that is, only a minor difference relative to what was found above when employing the repo rate.

²¹ It should be noted that the concept “zero lower bound” is something of a misnomer. In Sweden, the repo rate was lowered to -0.5% and other central banks have also had negative policy rates.

²² We use the shadow rate of De Rezende and Ristinieni (2018) based on a three-factor model. For other recent examples of analysis using shadow rates, see Wu and Xia (2016) and Conti (2017).

²³ Note that the data on the shadow rate range from 1996Q1 to 2018Q3. The model using these data is accordingly estimated on a marginally shorter sample than previous models.

Table 7 Log marginal likelihood and DIC for the different trivariate BVAR specifications using foreign variables. Price inflation and unemployment rate data

Model	US output gap		Oil prices	
	Log marginal likelihood	DIC	Log marginal likelihood	DIC
Constant	− 249.2	264.1	− 613.9	959.8
TVP	− 238.3	272.5	− 726.7	1207.8
SV	− 246.5	264.2	− 611.7	960.4
TVP-SV	− 235.9	274.9	− 722.4	1207.5

Based on trivariate BVARs with $y_t = (g_t^{\text{gapUS}} \ u_t \ \pi_t^p)'$ and $y_t = (\pi_t^{\text{oil}} \ u_t \ \pi_t^p)'$, where g_t^{gapUS} is the output gap in the USA, π_t^{oil} the year-on-year percentage change in the oil price, u_t is the unemployment rate and π_t^p is price inflation. Table gives the natural logarithm of the marginal likelihood of the models. “DIC” is the deviance information criterion. “Constant” is the BVAR model with constant covariance matrix and parameters. “TVP” is the BVAR model with constant covariance matrix and time-varying parameters. “SV” is the BVAR model with stochastic volatility and constant parameters. “TVP-SV” is the BVAR model with stochastic volatility and time-varying parameters

We next turn to the foreign variables. Through the inclusion of these, we aim to—at least to some extent—take into account the fact that Sweden is a small open economy and that it has been established that global factors have been shown to be empirically relevant determinants of inflation in several economies; see, for example, Conti et al. (2015) and Bobeica and Jarocinski (2017). We estimate two trivariate models—one where the US output gap is added to our benchmark bivariate system and the other where we instead add the percentage change (year-on-year) in the oil price. Data are shown in Fig. 25 in “Appendix”.²⁴

Results regarding model selection are shown in Table 7. Turning to the model including the US output gap first—that is, where we have set $y_t = (g_t^{\text{gapUS}} \ u_t \ \pi_t^p)'$ —it can be seen that the model with time-varying parameters and stochastic volatility is once again chosen when looking at marginal likelihoods.²⁵ The DIC on the other hand suggests that the model with constant parameters and covariance matrix is the preferred one. The difference to the “usual winner” when using the DIC—that is, the model with constant parameters and stochastic volatility—is extremely small though. Looking at the slope of the model with time-varying parameters and stochastic volatility in Fig. 26 in “Appendix”, we can note that it is very similar to what we found in our benchmark specification (shown in Fig. 5). The model with constant parameters and covariance matrix has a very moderate slope -0.02 , which can be compared to -0.06 which we found for the model with constant parameters and stochastic volatility in Sect. 3.

²⁴ The US output gap data range from 1995Q1 to 2018Q1. Similar to the the model using the shadow rate, it is accordingly estimated on a marginally shorter sample than the other models.

²⁵ We have chosen to place the US output gap first in the system. This reflects the fact that Sweden is a small open economy and accordingly should not affect a US variable. Obviously, the choice of ordering in the system does not alone make the US output gap exogenous with respect to the Swedish variables; Swedish variables can here affect the US output gap through lags. This could be changed, for example, through the use of additional hyperparameters which can forcefully shrink the coefficients of the Swedish variables in the equation for the US output gap to zero; see, for example, Villani and Warne (2003).

Finally, we note that for the model including oil price changes—that is, where $y_t = (\pi_t^{\text{oil}} u_t \pi_t^{\text{p}})'$ —model selection based on marginal likelihoods suggests that the model with constant parameters and stochastic volatility is the most appropriate. A model with constant parameters and covariance matrix is supported by the DIC. The estimated slope for these two models is -0.04 and -0.03 , respectively. The model selection hence points in a somewhat different direction than what we have seen for other specifications. Overall, it does not, however, change our main conclusion from the analysis conducted up until here: Even though the evidence in favour of a stable Phillips curve in Sweden is mixed, there is no evidence that the Swedish Phillips curve has become flatter in the last few years of the sample.

5 Out-of-sample forecasts

Above we have shown—perhaps somewhat ironically—that the mixed evidence in favour of a stable Phillips curve is a robust finding over many different specifications. The conclusion ultimately depends on which model selection tool one prefers. As a final empirical exercise, we will try to shed some light on the model selection issue by conducting an out-of-sample forecast exercise. There are admittedly pros and cons with this approach. In practice, we aim to establish whether there is Granger causality and Ashley et al. (1980) mean that out-of-sample forecasting is the natural way to do this. Diebold (2015), on the other hand, points out that any out-of-sample exercise makes inefficient use of the available information; while not completely dismissing out-of-sample forecasting performance as a tool for model selection, he concludes (p. 8) that “In general, full-sample model comparison procedures appear preferable”.

While we believe that Diebold has a valid point, we also believe that there might be additional information to be gained from an out-of-sample forecast exercise. First, there is a general risk that more flexible models are more prone to over-fitting the data and, as is widely known, over-fitting tends to hurt forecast performance.²⁶ In this case, our primary concern is that the model with both time-varying parameters and stochastic volatility—which is the most flexible specification—is most subject to this problem. Assessing this issue seems relevant. Second, and from a more general perspective, the number of out-of-sample forecast exercises conducted using BVARs with drifting parameters and/or stochastic volatility is fairly limited.²⁷ It therefore seems useful to also add to this literature.

²⁶ See, for example, Clark (2004) or Rossi and Sekhposyan (2011) for related discussions. It can also be noted that the traditional Minnesota prior employed in many Bayesian VARs—see, for example, Litterman (1986)—was developed to address the problem with over-fitting in VARs.

²⁷ We briefly mention a few important contributions here. Clark (2011) assessed the performance of a BVAR with stochastic volatility (and constant parameters and employing a steady-state prior) using US data on GDP growth, unemployment rate, inflation and the Federal funds rate. He concluded that there were modest gains to be made when it came to the accuracy of point forecasts but that relevant improvements could be made when it came to density forecasts. Using a BVAR with time-varying parameters and stochastic volatility, and employing US data on inflation, the unemployment rate and a short-term interest rate, D’Agostino et al. (2013) concluded that on average the model’s forecasting performance was better than that of the alternative models. The benefits from using the model with time-varying parameters and stochastic volatility were particularly large when forecasting inflation. Barnett et al. (2014) considered a fairly wide range of models when forecasting GDP growth, inflation and the 3-month treasury bill rate in the UK. One of their findings

The out-of-sample forecasting exercise is conducted using all four VAR models described earlier, that is, (i) constant covariance matrix and parameters, (ii) constant covariance matrix and time-varying parameters, (iii) stochastic volatility and constant parameters and (iv) stochastic volatility and time-varying parameters. In addition, we also introduce a univariate Bayesian AR model with constant parameters and covariance matrix for inflation—which is the variable we will focus upon—in order to get a simple benchmark model. The lag length is selected as five by the Akaike (1974) information criterion except for one forecast origin, 2005Q1, where the chosen lag length is one.²⁸ Regarding the variables used, it can be noted that we return to our main specification from Sect. 3 with the unemployment rate (age group 16–64 years) and CPIF inflation.

Out-of-sample forecasts are generated based on an expanding sample. We first estimate all models using data from 1995Q1 to 2004Q4 and then use the estimated models to generate forecasts up to eight quarters ahead with the median of the predictive distribution as the forecast. The sample is then extended one period, the models are re-estimated, and new forecasts up to eight quarters ahead generated. This continues until we generate the last forecasts based on the sample 1995Q1 to 2018Q2. We thus have a different number of forecasts to evaluate, from 55 forecasts at the one-quarter horizon to 48 at the eight-quarter horizon.

The results from the out-of-sample forecast exercise—in terms of the root mean square forecast errors (RMSFEs)—are given in Table 8. As can be seen, there is no single model that performs best at all horizons. The model with time-varying parameters and stochastic volatility is the best VAR model at most horizons (one to six quarters). It can be noted though that at the longest horizon, this model actually has the highest RMSFE of all models. At the two longest horizons (seven and eight quarters), it is instead the model with constant parameters and covariance matrix which performs best of the four VAR models. Differences are in general not particularly large though. For example, comparing the model with time-varying parameters and stochastic volatility to that with constant parameters and covariance matrix, the difference is never larger than 0.08 percentage points. Conducting a Diebold and Mariano (1995) test between these two models—under the assumption of quadratic loss—we cannot find statistical significance at any horizon.²⁹ Finding that the model with time-varying parameters and stochastic volatility has difficulties outperforming the model with constant parameters

Footnote 27 continued

was that a model with drifting parameters and stochastic volatility could generate improvements in forecast accuracy over a simple autoregressive benchmark (with constant parameters and variance) but that it was difficult to beat a VAR (with constant parameters and covariance matrix) estimated on a rolling sample. Aastveit et al. (2017) estimated models of various dimensions using US data but focused on the evaluation of the forecast accuracy with respect to GDP growth, unemployment rate and inflation. They found that the model with time-varying parameters and stochastic volatility had difficulties outperforming a model with constant parameters and covariance matrix.

²⁸ As an alternative, we have also used the Bayesian information criterion (Schwarz 1978) to select the lag length, with lag length one chosen for the first 42 forecasts (up to forecast origin 2015Q2) and five lags for the remaining forecast origins. This has only a marginal effect on the RMSEs of the forecasts, and the conclusions about the relative performance of the models are the same. Results are not reported but are available upon request.

²⁹ The absolute value of the test statistic is never larger than 0.98. Results are not reported but are available upon request. It can be noted that we use a “standard” Diebold and Mariano (1995) test even though the validity of this test in the present setting can be questioned. However, Diebold (2015) has pointed out that

Table 8 RMSFEs for inflation forecasts from Bayesian VARs and univariate Bayesian AR model from out-of-sample forecast exercise

Horizon	Constant	TVP	SV	TVP-SV	AR
1	0.386	0.379	0.383	0.365	0.364
2	0.527	0.485	0.518	0.482	0.487
3	0.614	0.561	0.603	0.545	0.563
4	0.710	0.652	0.701	0.631	0.657
5	0.730	0.660	0.725	0.643	0.674
6	0.725	0.693	0.723	0.685	0.679
7	0.712	0.717	0.718	0.718	0.685
8	0.698	0.765	0.705	0.769	0.689

Based on bivariate BVARs with $y_t = (u_t \pi_t^p)'$ where u_t is the unemployment rate and π_t^p is price inflation. “Constant” is the BVAR model with constant covariance matrix and parameters. “TVP” is the BVAR model with constant covariance matrix and time-varying parameters. “SV” is the BVAR model with stochastic volatility and constant parameters. “TVP-SV” is the BVAR model with stochastic volatility and time-varying parameters. “AR” is a univariate AR model for inflation. Out-of-sample forecasts generated using an expanding sample; the first sample is 1995Q1 to 2004Q4 and the last 1995Q1 to 2018Q2

and covariance matrix in out-of-sample forecasting, our results bear some resemblance to those of Aastveit et al. (2017).

It should, however, also be noted that the univariate AR model has the lowest RMSFE of all models at horizons one, six, seven and eight quarters. This finding is in line with what has been commonly pointed out in the literature on inflation forecasting, namely that simple univariate models tend to do quite well; see, for example, Stock and Watson (2009).

Given that the differences between the models are small, they should not be over-interpreted. What seems clear though is that the strong statements regarding the superiority of the model with time-varying parameters and stochastic volatility—for example, when it is compared to the model with constant parameters and covariance matrix—that the marginal likelihood comparisons in Table 1 show do not carry over to a clear improvement in out-of-sample forecast performance. One can also note regarding the VAR models that the model with constant parameters and stochastic volatility—which in previous sections received a lot of support from the DIC—never is the preferred specification according to the out-of-sample forecast analysis.

We conclude this section by noting that while marginal likelihood calculations tend to clearly favour the model with time-varying parameters and stochastic volatility, out-of-sample forecasts are less supportive of that model (even if it does perform best at several horizons). The fact that the differences in forecasting performance between the various models are small means that it is difficult to assess what the differences are due to. However, the finding that both the univariate AR model and the VAR with constant parameters and covariance matrix perform well in out-of-sample forecasting

Footnote 29 continued

if one uses Diebold and Mariano-type tests for model selection, one might as well use the “standard” test and associated critical values.

gives some support to old concern that in-sample measures might favour more complex models at the expense of forecasting performance.

6 Discussion

The analysis presented in Sects. 3 and 4 paints a fairly consistent picture. While the evidence is mixed regarding the stability of the Swedish Phillips curve, we consistently find that the Phillips curve has not been unusually flat during the last few years of the sample when inflation—despite efforts by the Riksbank to make monetary policy highly expansive—has been low. A flattened Phillips curve is, according to our analysis, hence not the cause for the low inflation in recent years.

Obviously, it is useful to be able to dismiss a potential explanation for the low inflation. But we would also of course like to know what caused it. The analysis presented in this paper suggests at least one explanation based on the model with time-varying parameters (and stochastic volatility). According to the model, the low inflation during the last few years of the sample can to some extent be explained by the low trend inflation that we found regardless of whether price or wage inflation was used. This might perhaps be perceived as a somewhat circular line of reasoning since the low outcomes for actual inflation are one reason for why the estimated trend inflation is low. But the findings of the model can be given an economic interpretation along the following lines: In the aftermath of the global financial crisis, the Riksbank has had substantial problems with making inflation reach the target rate. As can be seen from Fig. 1, CPIF inflation did not reach 2% for 6 years in a row, 2011–2016. In light of this development, economic agents in Sweden began to lose confidence in the inflation target which manifested itself in falling inflation expectations among other things. As pointed out above, this was a main reason for the Riksbank to abandon its leaning against the wind policy in 2014. The model's low trend inflation can be seen as reflecting the low inflation expectations that we saw during this period. There is evidence that the low inflation outcomes affected the wage formation process; as employers and employees no longer expected inflation around the target, wage agreements with low nominal increases were signed. Seeing that wage agreements tend to run for several years, this means that such effects based on expectations will be fairly long-lasting. As price inflation has increased, so has trend inflation in the model relying on prices. As new wage agreements will be negotiated in a period with higher price inflation, it is reasonable to assume that wage increase claims will be higher and that both wage inflation and the accompanying measure of trend inflation will increase.³⁰

³⁰ That higher wage inflation will drive up price inflation is in line with our theoretical expectations as discussed above. It also finds support in modelling work that we have conducted. Estimating yet another bivariate BVAR with time-varying parameters and stochastic volatility, now with wage and price inflation as the variables in the system—that is, we set $y_t = (\pi_t^w, \pi_t^p)'$ —we find that a shock to wage inflation increases price inflation; see Fig. 27 in “Appendix”. But while this empirical evidence points in the right direction, it can also be noted that the variance decomposition from this model with wage and price inflation indicates that shocks to wage inflation explain little of the forecast error variance of price inflation; see Fig. 28 in “Appendix”. These results are in line with the findings of Peneva and Rudd (2017) who established that changes in labour costs seemed to have little effect on price inflation in the USA.

Summing up, we argue that part of the problem with low inflation that Sweden has experienced recently may have been self-inflicted and a consequence of the Riksbank's policy of leaning against the wind. That said, we also want to point out that other factors of course have affected inflation during this period. It can be noted that the broad findings put forward in this paper are in line with those of Karlsson and Österholm (2018) regarding the US Phillips curve; together with the fact that inflation developments over the last few years have been similar in several countries, this indicates that there are international determinants of inflation at work as well, for instance the commonly suggested globalization and digitalization.

7 Conclusions

In this paper we have assessed the stability of the Swedish Phillips curve. Conducting model selection using new methods, we found mixed evidence in favour a stable relationship between inflation and the unemployment rate. Based on this finding, we conclude that it is unclear whether the Phillips curve has been stable over time. However, while support for time-varying parameters (and stochastic volatility) can be found when looking at the models' marginal likelihood, our results also clearly indicate that the low inflation in Sweden in recent years is not due to a flatter Phillips curve.

In terms of explaining these low inflation outcomes in recent years, the models with time-varying parameters suggest that low trend inflation—reflecting low inflation expectations which likely affected the wage formation process—has contributed. This is in line with the findings of Clark (2014) and Watson (2014) who both point to the importance of a well-anchored trend level of inflation when it comes to explaining inflation outcomes in the USA. Other factors, such as globalisation and digitalisation, are likely to be important parts of the explanation as well given that many other countries also have seen low inflation outcomes during the same period. The results presented in this paper thereby add to our knowledge both concerning the properties of the Swedish Phillips curve but also international developments, although there is admittedly further work needed. From a more methodological viewpoint, it is worth pointing out that such further work should—to a larger extent than what is being done currently—consider the possibility of relying on frameworks that allow for time variation in the relationships between macroeconomic variables as well as the variance of the shocks that hit the economy.

Acknowledgements Open access funding provided by Örebro University.

Compliance with ethical standards

Conflict of interest None of the authors has any conflicts of interest.

Ethical approval This article does not contain any studies with human participants or animals performed by any of the authors.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

Appendix

See Figs. 11, 12, 13, 14, 15, 16, 17, 18, 19, 20, 21, 22, 23, 24, 25, 26, 27 and 28.

Fig. 11 Impulse-response function: Effect of shocks to the unemployment rate on the unemployment rate. Price inflation data, model with constant parameters and stochastic volatility. *Note:* Based on bivariate BVAR with $y_t = (u_t \pi_t^p)'$, where u_t is the unemployment rate and π_t^p is price inflation. Size of impulse is one standard deviation. Effect in percentage points on vertical axis. Horizon in quarters and dates on horizontal axes

Fig. 12 Impulse-response function: Effect of shocks to inflation on inflation. Price inflation data, model with constant parameters and stochastic volatility. *Note:* Based on bivariate BVAR with $y_t = (u_t \pi_t^p)'$, where u_t is the unemployment rate and π_t^p is price inflation. Size of impulse is one standard deviation. Effect in percentage points on vertical axis. Horizon in quarters and dates on horizontal axes

Fig. 13 Impulse-response function: Effect of shocks to inflation on the unemployment rate. Price inflation data, model with constant parameters and stochastic volatility. *Note:* Based on bivariate BVAR with $y_t = (u_t \pi_t^p)'$, where u_t is the unemployment rate and π_t^p is price inflation. Size of impulse is one standard deviation. Effect in percentage points on vertical axis. Horizon in quarters and dates on horizontal axes

Fig. 14 Impulse-response function: Effect of shocks to the unemployment rate on the unemployment rate. Price inflation data, model with time-varying parameters and stochastic volatility. *Note:* Based on bivariate BVAR with $y_t = (u_t \pi_t^p)'$, where u_t is the unemployment rate and π_t^p is price inflation. Size of impulse is one standard deviation. Effect in percentage points on vertical axis. Horizon in quarters and dates on horizontal axes

Fig. 15 Impulse-response function: Effect of shocks to inflation on inflation. Price inflation data, model with time-varying parameters and stochastic volatility. *Note:* Based on bivariate BVAR with $y_t = (u_t \pi_t^p)'$, where u_t is the unemployment rate and π_t^p is price inflation. Size of impulse is one standard deviation. Effect in percentage points on vertical axis. Horizon in quarters and dates on horizontal axes

Fig. 16 Impulse-response function: Effect of shocks to inflation on the unemployment rate. Price inflation data, model with time-varying parameters and stochastic volatility. *Note:* Based on bivariate BVAR with $y_t = (u_t \pi_t^p)'$, where u_t is the unemployment rate and π_t^p is price inflation. Size of impulse is one standard deviation. Effect in percentage points on vertical axis. Horizon in quarters and dates on horizontal axes

Fig. 17 Estimated coefficients of the unemployment equation from the BVAR in Eq. (13). Price inflation and unemployment rate data. Model with time-varying parameters and stochastic volatility. *Note:* Based on bivariate BVARs with $y_t = (u_t \pi_t^p)'$, where u_t is the unemployment rate and π_t^p is price inflation

Fig. 18 Impulse-response function: Effect of shocks to the unemployment rate on wage inflation. Model with constant parameters and stochastic volatility. *Note:* Based on bivariate BVAR with $\mathbf{y}_t = (u_t \pi_t^w)'$ where u_t is the unemployment rate and π_t^w is wage inflation. Size of impulse is one standard deviation. Effect in percentage points on vertical axis. Horizon in quarters and dates on horizontal axes

Fig. 19 Equilibrium unemployment rate, unemployment gap and output gap. *Note:* Equilibrium unemployment rate measured in percent. Unemployment gap measured in percentage points. Output gap measured in percent of potential GDP. *Source:* National Institute of Economic Research

Fig. 20 Estimated “slope” of the Phillips curve. Price inflation and unemployment gap data. Model with time-varying parameters and stochastic volatility. *Note:* Based on bivariate BVAR with $y_t = (u_t^{\text{gap}} \pi_t^{\text{p}})'$, where u_t^{gap} is the unemployment gap and π_t^{p} is price inflation. Coloured band is 68% equal-tail credible interval

Fig. 21 Estimated “slope” of the Phillips curve. Price inflation and output gap data. Model with time-varying parameters and stochastic volatility. *Note:* Based on bivariate BVAR with $y_t = (g_t^{\text{gap}} \pi_t^{\text{p}})'$, where g_t^{gap} is the output gap and π_t^{p} is price inflation. Coloured band is 68% equal-tail credible interval

Fig. 22 Repo rate and shadow rate. *Note:* Both variables are measured in percent. *Source:* National Institute of Economic Research and De Rezende and Ristinieni (2018)

Fig. 23 Estimated “slope” of the Phillips curve from trivariate BVAR including repo rate. Price inflation and unemployment rate data. Model with time-varying parameters and stochastic volatility. *Note:* Based on trivariate BVAR with $y_t = (u_t \pi_t^p i_t^r)'$. Coloured band is 68% equal-tail credible interval

Fig. 24 Estimated “slope” of the Phillips curve from trivariate BVAR including shadow rate. Price inflation and unemployment rate data. Model with time-varying parameters and stochastic volatility. *Note:* Based on trivariate BVAR with $y_t = (u_t \pi_t^p i_t^s)'$. Coloured band is 68% equal-tail credible interval

Fig. 25 Output gap in the USA and change in oil price. *Note:* Output gap measured in percentage points. Change in oil price is the year-on-year change given in percent. *Source:* FRED

Fig. 26 Estimated “slope” of the Phillips curve from trivariate BVAR including US output gap. Price inflation and unemployment rate data. Model with time-varying parameters and stochastic volatility. *Note:* Based on trivariate BVAR with $y_t = \begin{pmatrix} g_t^{\text{gapUS}} \\ u_t \\ \pi_t^p \end{pmatrix}'$, where g_t^{gapUS} is the output gap in the USA, u_t is the unemployment rate and π_t^p is price inflation. Coloured band is 68% equal-tail credible interval

Fig. 27 Impulse-response function: Effect of shocks to wage inflation on price inflation. Model with constant parameters and stochastic volatility. *Note:* Based on bivariate BVAR with $y_t = (\pi_t^w \pi_t^p)'$, where π_t^w is wage inflation and π_t^p is price inflation. Size of impulse is one standard deviation. Effect in percentage points on vertical axis. Horizon in quarters and dates on horizontal axes

Fig. 28 Variance decomposition: Share of forecast error variance of price inflation explained by shocks to wage inflation. Model with time-varying parameters and stochastic volatility. *Note:* Based on bivariate BVAR with $y_t = (\pi_t^w \pi_t^p)'$, where π_t^w is wage inflation and π_t^p is price inflation. Share on vertical axis. Horizon in quarters and dates on horizontal axes

References

- Aastveit KA, Carriero A, Clark TE, Marcellino M (2017) Have standard VARs remained stable since the crisis? *J Appl Econ* 32:931–951
- Akaike H (1974) A new look at the statistical model identification. *IEEE Trans Autom Control* 19:716–723
- Akram F, Mumtaz H (2019) Time-varying dynamics of the Norwegian economy. *Scand J Econ* 121:407–434
- Ashley R, Granger CWJ, Schmalensee R (1980) Advertising and aggregate consumption: an analysis of causality. *Econometrica* 48:1149–1167
- Ball L, Mazumder S (2011) Inflation dynamics and the great recession. *Brook Pap Econ Activity* 42:337–405
- Barnett A, Mumtaz H, Theodoridis K (2014) Forecasting UK GDP growth and inflation under structural change. a comparison of models with time-varying parameters. *Int J Forecast* 30:129–143
- Bell DNF, Blanchflower DG (2018) Underemployment in the US and Europe. NBER Working Paper No. 24927
- Benati L, Surico P (2008) Evolving U.S. monetary policy and the decline of inflation predictability. *J Eur Econ Assoc* 6:634–646
- Bernanke B (1986) Alternative explanations of the money-income correlation. *Carnegie-Rochester Conf Ser Public Policy* 25:49–100
- Blanchard OJ, Quah D (1989) The dynamic effects of aggregate demand and supply disturbances. *Am Econ Rev* 79:655–673
- Blanchard O, Cerutti E, Summers L (2015) Inflation and activity—two explorations and their monetary policy implications. NBER Working Paper No. 21726
- Bobeica E, Jarocinski M (2017) Missing disinflation and missing inflation: the puzzles that aren't. Working Paper No. 2000, European Central Bank
- Bulligan G, Guglielminetti E, Viviano E (2017) Wage growth in the Euro area: where do we stand? *Occasional Papers* 413, Bank of Italy
- Busetti F, Caivano M (2016) The trend-cycle decomposition of output and the Phillips Curve: Bayesian estimates for Italy and the Euro area. *Empir Econ* 50:1565–1587
- Chan JCC, Eisenstat E (2018) Bayesian model comparison for time-varying parameter VARs with stochastic volatility. *J Appl Econ* 33:509–532
- Chan JCC, Koop G, Potter SM (2016) A bounded model of time variation in trend inflation, Nairu and the Phillips curve. *J Appl Econ* 31:551–565
- Clark TE (2004) Can out-of-sample forecast comparisons help prevent overfitting? *J Forecast* 23:115–139
- Clark TE (2011) Real-time density forecasts from Bayesian vector autoregressions with stochastic volatility. *J Bus Econ Stat* 29:327–341
- Clark TE (2014) The importance of trend inflation in the search for missing disinflation. *Federal Reserve Bank of Cleveland Economic Commentary* 2014–2016
- Clark TE, Doh T (2014) Evaluating alternative models of trend inflation. *Int J Forecast* 30:426–448
- Clark TE, McCracken MW (2006) The predictive content of the output gap for inflation: resolving in-sample and out-of-sample evidence. *J Money Credit Bank* 38:1127–1148
- Cogley T, Sargent TJ (2005) Drifts and volatilities: monetary policies and outcomes in the post WWII US. *Rev Econ Dyn* 8:262–302
- Cogley T, Sbordone AM (2009) Trend inflation, indexation, and inflation persistence in the new Keynesian Phillips curve. *Am Econ Rev* 98:2101–2126
- Cogley T, Primiceri GE, Sargent TJ (2009) Inflation-Gap Persistence in the US. *Am Econ J Macroecon* 2:43–69
- Coibion O, Gorodnichenko Y (2015) Is the Phillips curve alive and well after all? Inflation expectations and the missing disinflation. *Am Econ J Macroecon* 7:197–232
- Conti AM (2017) Has the FED fallen behind the curve? Evidence from VAR models. *Econ Lett* 159:164–168
- Conti AM, Gigante C (2018) Weakness in Italy's core inflation and the Phillips Curve: the role of labour and financial indicators. *Occasional Papers* 466, Bank of Italy
- Conti AM, Neri S, Nobili A (2015) Why is inflation so low in the Euro area? *Economic Working Papers* 1019, Bank of Italy
- Conti AM, Neri S, Nobili A (2017) Low inflation and monetary policy in the Euro area. *ECB Working Paper No.* 2005
- D'Agostino A, Gambetti L, Giannone D (2013) Macroeconomic forecasting and structural change. *J Appl Econ* 28:82–101

- De Rezende RB, Ristiniemi A (2018) A shadow rate without a lower bound constraint. Sveriges Riksbank Working Paper Series No. 355
- Diebold FX (2015) Comparing predictive accuracy, twenty years later: a personal perspective on the use and abuse of Diebold–Mariano tests. *J Bus Econ Stat* 33:1–9
- Diebold FX, Mariano RS (1995) Comparing predictive accuracy. *J Bus Econ Stat* 13:253–263
- Faust J, Wright JH (2013) Forecasting inflation. In: Elliott G, Granger CWJ, Timmermann A (eds) *Handbook of economic forecasting*, vol 2A. North-Holland, Amsterdam
- Franta M, Horvath R, Rusnak M (2014) Evaluating changes in the monetary transmission mechanism in the Czech Republic. *Empir Econ* 46:827–842
- Galí J, Gertler M (1999) Inflation dynamics: a structural econometric analysis. *J Monet Econ* 44:195–222
- Hall RE (2003) Modern theory of unemployment fluctuations: empirics and policy applications. *Am Econ Rev* 93:145–150
- Huang Y-F, Luo S (2018) Potential output and inflation dynamics after the great recession. *Empir Econ* 55:495–517
- IMF (2013) *World Economic Outlook*, April 2013
- Jansson P (2017) The ideological debate on monetary policy—lessons from developments in Sweden. Speech given at Fores, Stockholm, December 6
- Johansson J, Löf M, Sigrist O, Tysklind O (2018) Measures of core inflation in Sweden. *Economic Commentaries* No. 11/2018, Sveriges Riksbank
- Juselius M, Takáts E (2018) The enduring link between demography and inflation. *BIS Working Papers* No. 722
- Karlsson S, Österholm P (2018) Is the US Phillips curve stable? Evidence from Bayesian VARs. Working Paper 2018:5, School of Business, Örebro University
- Kim S, Shephard N, Chib S (1998) Stochastic volatility: likelihood inference and comparison with ARCH models. *Rev Econ Stud* 65:361–393
- King RG, Watson MW (1994) The post-War U.S. Phillips curve: a revisionist econometric history. *Carnegie-Rochester Conf Ser Public Policy* 41:157–219
- Knotek ES, Zaman S (2017) Have inflation dynamics changed? Federal Reserve Bank of Cleveland *Economic Commentary* 2017–2021
- Koop G, Leon-Gonzalez R, Strachan RW (2009) On the evolution of the monetary policy transmission mechanism. *J Econ Dyn Control* 33:997–1017
- Leduc S, Wilson DJ (2017) Has the Wage Phillips curve gone Dormant? Federal Reserve Bank of San Francisco *Economic Letter* 2017–2030
- Litterman RB (1986) Forecasting with Bayesian vector autoregressions: five years of experience. *J Bus Econ Stat* 5:25–38
- Matheson T, Stavrev E (2013) The Great recession and the inflation puzzle. *Econ Lett* 120:468–472
- National Institute of Economic Research (2017) *Wage Formation in Sweden 2017*
- Peneva EV, Rudd J (2017) The passthrough of labor costs to price inflation. *J Money Credit Bank* 49:1777–1802
- Phillips AW (1958) The Relationship between unemployment and the rate of change of money wages in the United Kingdom 1861–1957. *Economica* 25:283–299
- Primiceri G (2005) Time varying structural vector autoregressions and monetary policy. *Rev Econ Stud* 72:821–852
- Riggi M, Venditti F (2015) Failing to forecast low inflation and Phillips curve instability: a Euro-area perspective. *Int Finance* 18:47–68
- Rossi B, Sekhposyan T (2011) Understanding models' forecasting performance. *J Econom* 164:158–172
- Rudd J, Whelan K (2005) New tests of the new-Keynesian Phillips curve. *J Monet Econ* 52:1167–1181
- Schwarz G (1978) Estimating the dimension of a model. *Ann Stat* 6:461–464
- Spiegelhalter DJ, Best NG, Carlin BO, van der Linde A (2002) Bayesian measures of model complexity and fit. *J R Stat Soc B* 64:583–639
- Stock JH, Watson MW (2009) Phillips curve inflation forecasts. In: Fuhrer J, Kodrzycki Y, Little J, Olivei G (eds) *Understanding inflation and the implications for monetary policy*. MIT Press, Cambridge
- Svensson LEO (2014) Inflation targeting and “leaning against the wind”. *Int J Central Bank* 10:103–114
- Svensson LEO (2015) The possible unemployment cost of average inflation below a credible target. *Am Econ J Macroecon* 7:258–296
- Sveriges Riksbank (1993) The Riksbank's target for monetary policy. Press Release No. 5, January

- Villani M, Warne A (2003) Monetary policy analysis in a small open economy using Bayesian cointegrated structural VARs. Working Paper No. 296, European Central Bank
- Watson MW (2014) Inflation persistence, the NAIRU, and the great recession. *Am Econ Rev* 104:31–36
- Williamson S (2018) Inflation control: do central bankers have it right? *Fed Reserve Bank St Louis Rev* 100:127–150
- Wu JC, Xia FD (2016) Measuring the macroeconomic impact of monetary policy at the zero lower bound. *J Money Credit Bank* 48:253–291
- Yellen J (2017) Transcript of Chair Yellen's Press conference. Washington DC, September 20

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.